# Ozy: A General Orchestration Container

Glenn Osborne
*Catavolt Inc*
*Alpharetta, GA 30022, USA*
*Email: glenn.osborne@catavolt.com*

Tim Weninger
*Department of Computer Science and Engineering*
*University of Notre Dame*
*Notre Dame, IN 46556*
*Email: tweninge@nd.edu*

arXiv:1604.07642v1 [cs.SE] 25 Apr 2016

***Abstract*—Service-Oriented Computing is a paradigm that uses services as building blocks for building distributed applications. The primary motivation for orchestrating services in the cloud used to be distributed business processes, which drove the standardization of the Business Process Execution Language (BPEL) and its central notion that a service is a business process. In recent years, there has been a transition towards other motivations for orchestrating services in the cloud, *e.g.*, XaaS, RMAD. Although it is theoretically possible to make all of those services into WSDL/SOAP services, it would be too complicated and costly for industry adoption. Therefore, the central notion that a service is a business process is too restrictive. Instead, we view a service as a technology neutral, loosely coupled, location transparent procedure. With these ideas in mind, we introduce a new approach to services orchestration: Ozy, a general orchestration container. We define this new approach in terms of existing technology, and we show that the Ozy container relaxes many traditional constraints and allows for simpler, more feature-rich applications.**

***Keywords*-SOA, SOC, coordination, orchestration, service**

## I. Introduction

Service-Oriented Computing (SOC) is a paradigm that uses services as building blocks for building distributed applications. Services perform procedures, which can be anything from simple information requests to complicated business processes. The original focus of SOC was the automation of business processes by coordinating the services of multiple business partners. In recent years, however, the Web has evolved into an ecosystem of Everything-as-a-Service (XaaS) [1]; consequently, services orchestration has moved beyond the need for just coordinating business processes and has evolved into a more general problem of programming-in-the-large.

The basic Service-Oriented Architecture (SOA) is a relationship involving three kinds of participants: Service Provider, Service Client and Service Registry. The Extended Service-Oriented Architecture (ESOA) extends the basic SOA with three tiers that address overarching concerns. The middle tier is a service composition layer that facilitates orchestrations over service definitions and service implementations [2].

The standard language for service orchestration is the Business Process Execution Language (WS-BPEL or BPEL), which is based on the specifications for SOAP, WSDL and UDDI [3]. Its design is centered on the notion of the business process being the glue between interacting services.

In recent years, RESTful (REpresentational State Transfer) services have emerged as a lightweight alternative to SOAP services. Whereas SOAP services are based on remote procedure calls, RESTful services build upon the HTTP hypermedia standards of resources, resource identifiers and representations. The key elements of RESTful services are technology neutral components, connectors and resource identifiers [4].

The standard technique for composing RESTful Web services is the ”mashup”. There are roughly three kinds of mashups: data-oriented, process-oriented and consumer-oriented. Unlike data and consumer-oriented mashups, which permeate the Web, process-oriented mashups are less frequent, but most typically provide for the composition of business processes [5].

There are several forces driving the proliferation and heterogeneity of service-oriented computing. The adoption of utility computing on a pay-as-you-go basis has allowed enterprises to develop applications more easily, cost effectively, and reliably than ever before [6]. New service platforms, such as Rapid Mobile App Development (RMAD) [7], facilitate opportunistic software development by business analysts. Finally, the need to modernize legacy systems has led to the micro-container concept that uses services to provide accessibility and increased scalability of older software architectures [8].

The Web service concept has evolved and now means more than just a business process. As a result, we believe that the industry can benefit from a more general treatment of services orchestration. Therefore, in this paper we introduce a new approach to services orchestration with Ozy, a general orchestration container based on the Oz Computation Model (OCM) [9]. In addition to the description provided in the present work we have released the orchestration framework as an open source software project under the Affero GPL at https://github.com/ozyio.

To illustrate this new approach, its benefits and its restrictions, we show examples of an RMAD platform for both the development and deployment of mobile apps. As illustrated in Fig. 1, large enterprises typically orchestrate several kinds

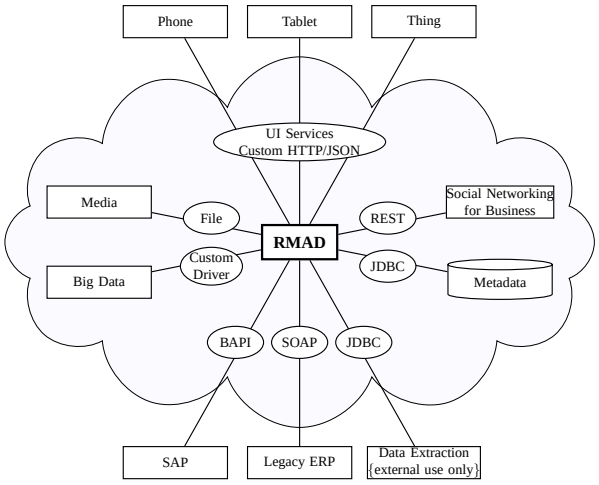


Figure 1. Heterogeneous services on an enterprise system. The Ozy container aims to elegantly orchestrate these diverse services.

of services including:

- Swagger/RESTful services
- WSDL/SOAP service
- Database services (NOSQL and JDBC)
- File services for serving media data
- Legacy services to access SAP (BAPI) or other older ERP services
- Proprietary HTTP/JSON services to support rich user experiences [10]

In its original sense, a service is simply a technology neutral, loosely coupled, location transparent procedure [2]. However, the prevailing techniques for orchestrating services are lacking support for the heterogeneity illustrated here. BPEL, for example, has demonstrated success in orchestrating business processes, but its central notion that a service is a business process is too restrictive. Furthermore, BPEL suffers from two strong criticisms: (1) it lacks a formal specification, and (2) its language is based on a cumbersome XML-based syntax [11]. RESTful orchestration techniques are also too restrictive because they also focus on business processes [12]. Moreover, the RESTful orchestration techniques that do exist are immature [5].

In the present work we present Ozy, a general orchestration container. Ozy is based on the Oz Computation Model (OCM); although the Oz language is best known for teaching multi-paradigm programming practices [13], the underlying theory and constructs of the OCM are especially useful to services orchestration. Specifically, we exploit implicit synchronization, dataflow variables [14], lazy execution, failed values and partial termination [9]. Furthermore, the expressibility of the Oz language allows us to use programming techniques not typically found in orchestration languages, such as declarative functional programming and pattern matching [15]. To the best of our knowledge, Ozy is the only orchestration engine that supports both implicit synchronization and persistent execution state. In addition, the Ozy orchestration container also includes the following contributions:

1) A new notion of partial activation as a dual to the existing notion of partial termination
2) A technology neutral orchestration architecture inspired by the elements of a network-based architecture
3) An open source orchestration container that implements the orchestration architecture (2), implements an Oz language interpreter based on the OCM, and exploits the notions of partial activation (1) and partial termination to support a persistent execution state

This paper describes the Ozy framework with special attention paid to the limitations of current approaches. After a section on related work and a brief introduction to the Oz computation model, the Ozy architecture and container systems are described. We finally provide two illustrative examples that highlights the Ozy orchestration container in real world scenarios.

## II. Related Work

Service-Oriented Computing (SOC) is a paradigm that uses loosely-coupled services as the fundamental elements for building applications. SOC relies on the Services-Oriented Architecture (SOA) and the Extended Services-Oriented Architecture (ESOA) to define standards for publishing, finding, binding and composing services [2]. Service-Oriented Middleware (SOM) supports the service-oriented interaction pattern with the provisioning of software features for deploying, publishing, finding, and binding services at runtime [16], [17].

Composing services is an exercise in programming-in-the-large. DeRemer and Kron argue that programming-in-the-large is a distinct and different intellectual activity from programming-in-the-small, and consequently, a different ”module interconnect language” should be used instead of a common programming language, such as one typically used to implement modules.

In the traditional view, there are three service composition models: Choreography, Orchestration and Wiring. Orchestration is the most used composition model, and the de facto standard is BPEL, an orchestration language based on the base specifications of SOAP, WSDL and UDDI [3]. Although BPEL is the industry standard, it suffers from two big criticisms: (1) it lacks formalization and (2) is based on a cumbersome XML syntax [11].

SOCK and JOLIE are two more-recent contributions that address some of the deficiencies of BPEL. SOCK (Service-Oriented Computing Kernel) formalizes the basic mechanisms of services communications and composition, therefore addressing the lack of formalism in BPEL [18]. JOLIE (Java Orchestration Language Interpreter Engine) builds on SOCK to create an expressive language and

orchestration interpreter engine, therefore addressing the cumbersome XML-based syntax of BPEL [19], [20].

The open source Apache ODE (Orchestration Director Engine) project is the underlying execution engine for several BPEL implementations, such as JBoss RiftSaw and WSO2 Business Process Server. Apache ODE relies on the Apache JaCOb (Java Concurrent Objects) framework to implement two key features required of BPEL processes: concurrency and persistence of execution state.

In recent years, and because of the success of the RESTful architecture, a Web service is more generally defined as a SOAP-based or RESTful-based service that can be described, advertised and discovered through Internet-based protocols [21]. SOAP services are based on remote procedure calls. In contrast, RESTful services build on the HTTP hypermedia standards of resources, resource identifiers and representations. The key elements of RESTful services are technology neutral components, connectors and resource identifiers [4].

Lightweight RESTful services are designed to ease consumption, composition and have been successfully used to build community-driven services known as mashups [22], [5]. Another approach for composing RESTful services is a technique that uses BPEL to orchestrate RESTful services by wrapping RESTful services with WSDL definitions [12].

Cloud computing has emerged in conjunction with Web services as a dominant utility computing solution. Cloud vendors offer platform-as-a-service (PaaS) and software-as-a-service (SaaS) applications on a pay-as-you-go basis allowing enterprises to develop applications more easily, cost effectively, and reliably than ever before [6]. Cloud services have become so prolific that we now refer to the provider ecosystem as Everything-as-a-Service (XaaS) [1].

Opportunistic software methods are being used in cloud computing to drive down the backlog of applications by allowing non-traditional software developers to combine two or more software systems to create functionality that mainstream software does not provide [23]. For example, Gartner recently identified the Rapid Mobile App Development (RMAD) category of development tools [7]. RMAD is opportunistic software development for mobile apps that offers simpler, faster development by business analysts.

The need to reduce application backlogs is just one of many motivating forces affecting service-oriented computing in the cloud. Service-oriented computing has focused primarily on easy integration at the server-side at the price of lower flexibility and ease-of-use at the client-side [24]. Focus is shifting to client-side interactions with services that sometimes require innovative service protocols [10]. Another force affecting cloud computing is the need to modernize legacy software. The micro-container concept uses services to provide accessibility and increased scalability of older software architectures [8].

Table I
FEATURE COMPARISON OF EXISTING ORCHESTRATION LANGUAGES: BPEL, JOLIE AND THE PRESENT WORK OZY. OZ, ITSELF, IS NOT AN ORCHESTRATION LANGUAGE, BUT IS INCLUDED HERE FOR COMPARISON.

| | BPEL | JOLIE | Oz | **Ozy** |
|---|---|---|---|---|
| Formal Specification | - | + | + | + |
| High-Level Language | - | + | + | + |
| Interpreted | + | + | - | + |
| Functional Programming | - | - | + | + |
| Long Running | + | + | - | + |
| Declarative Concurrency | + | + | + | + |
| Implicit Synchronization | - | - | + | + |
| Pattern Matching | - | - | + | + |
| Unification | - | - | + | + |
| Entailment | - | - | + | + |

*The Missing Link*

Composing services, and orchestrating in particular, has moved beyond the need for coordinating business processes and has become a more general programming-in-the-large problem. Although the Oz programming language is best known for its support of multi-paradigm programming [13], the theory and constructs of the Oz Computation Model (OCM) are especially useful for services orchestration. Most notable are implicit synchronization, dataflow variables [14], lazy execution, failed values and partial termination [9]. Furthermore, the expressibility of the Oz language enables programming techniques not typically found in programming-in-the-large [15]. In the present work, we adapt the Oz Computational Model for practical use in service orchestration. We call this adaption Ozy, and show that it is a more elegant, feature rich orchestration container.

## III. THE OZ COMPUTATION MODEL

The Oz Computation Model (OCM) [9] contains abstractions especially useful for services orchestration. Specifically, we exploit *implicit synchronization*, *dataflow variables* [14], *lazy execution*, *failed values* and *partial termination* [9] provided by the OCM. Furthermore, the expressibility of the Oz language allows advanced programming techniques, such as declarative functional programming and pattern matching [15].

The *dataflow variable* is the most powerful OCM abstraction because it greatly simplifies concurrency by hiding explicit synchronization, thereby giving us *implicit synchronization*. Dataflow variables provide implicit synchronization at the most atomic level–a single variable. A dataflow variable is created and bound in separate steps. The single rule of accessing a dataflow variable is very simple: *execution waits until the dataflow variable is bound. Once bound, execution continues*. This property allows advanced programming techniques never before seen in an orchestration language. The power of dataflow variables is best explained by Smolka:

> A basic problem with existing programming languages is that they delegate the creation and coordination of

concurrent computational activities to the underlying operating system and network protocols. This has the severe disadvantage that the *data abstractions of the programming language cannot be shared between communicating computational agents*. Thus the benefits of existing programming languages do not extend to the central concerns of concurrent and distributed software systems [14] (emphasis added).

The present work, Ozy, uses the dataflow variable abstraction thereby allowing service messages to be shared among interacting services independent of the underlying operating system and network protocols.

For the following examples on dataflow variables, the complete Oz syntax can be found in [9]. Our simple example, illustrated in Fig. 2, executes 3 concurrent threads. The third thread waits on variables $X$ and $Y$ before continuing. This is an example of implicit synchronization because the third thread is void of any syntax that explicitly waits for variables $X$ and $Y$ to be bound.

```
1  thread X = 5 end
2  thread Y = 7 end
3  thread Z = X + Y end
```

Figure 2. Simple example of implicit synchronization with 3 threads. Despite the lack of explicit syntax, implicit synchronization via dataflow variables requires that the third thread waits for variables $X$ and $Y$ to be bound (*i.e.*, thread 1 and 2 to complete) before continuing.

A more sophisticated example of dataflow programming illustrated in Fig. 3 involves nested data structures and function calls[1]. Dataflow variables can be arbitrarily nested in both data structures and logic. Because synchronization is implicit, logic can be written very cleanly, void of explicit techniques that synchronize concurrency.

Although our sophisticated example appears intuitive, there are actually deep dependencies on concurrency and synchronization: (1) two numeric values are compared that (2) depend on the result of pattern matching that (3) depends on the result of a function call to `GetCustomerInfo` that constructs a compound value that depends on two other nested function calls to `GetAverageSale` and `GetLastSale` where (4) one of the nested function calls depends on (5) an answer that (6) is fulfilled concurrently and takes three days to complete. Dataflow variables greatly simplify concurrent programming. Without dataflow variables, this example written in a mainstream programming language would require explicit synchronization or a concurrency library, such as monadic combinators for a functional language or an actor library for an imperative language.

Another key OCM abstraction is the Kernel Language, a concise set of values and statements that serve as the

[1]In Oz, function signatures and function applications are delimited with curly braces.

```
1  fun {GetCustomerInfo Id}
2    fun {GetAverageSale Id} Answer in
3      thread                                  <- 6
4        {Sleep 3 days}
5        Answer = 12345.67
6      end
7      return Answer                           <- 5
8    end
9    fun {GetLastSale Id}
10     123                                     <- 3
11   end
12   customer(key:Id averageSale:{GetAverageSale Id} ↩
         lastSale:{GetLastSale Id})            <- 4
13 end
14 local C D in
15   C = {GetCustomerInfo 'acme'}              <- 2
16   case C of customer(key:Id averageSale:AS lastSale:LS)
17     D = if LS > AS then 'Y' else 'N' end     <- 1
18   end
19 end
```

Figure 3. Sophisticated example of dataflow programming that computes a discount $D$ based on a complex, concurrent and implicitly synchronized work flow. Existing orchestration languages are ill-equipped to elegantly orchestrate this process.

basic computational elements and as building blocks for higher level computing techniques. The OCM defines formal operational semantics and an abstract machine that specify exactly how to execute a kernel language program. Informally, the abstract machine is a process consisting of several semantic stacks and a shared store of dataflow variables. Each semantic stack corresponds to a thread, and all semantic stacks access the same store. Execution begins with an initial process state and proceeds as a sequence of state transitions where a semantic stack (thread) is chosen and reduced. Processing continues until all semantic stacks are null. Semantic stacks can be reduced in parallel with synchronization on the store of dataflow variables.

More formally, the Oz computation model is based on the following definitions:

$\sigma$ = single-assignment store (dataflow variables)
$s$ = kernel statement
$E$ = environment
$(s, E)$ = semantic statement
$[(s, E)]$ = semantic stack (ST)
$\{[(s, E)]\}$ = multiset semantic stack (MST)

A computation step begins by choosing from the multiset one runnable (reducable) semantic stack giving us the following definition:

(a) $\text{MST} = \{\text{ST}\} \uplus \text{MST}'$

With a semantic stack chosen, a computation step is completed with the following definition:

(b) $(\{\text{ST}\} \uplus \text{MST}', \sigma) \rightarrow (\{\text{ST}'\} \uplus \text{MST}', \sigma')$

Program execution repeats steps (a) and (b) until termination [9].

The OCM executes over dataflow variables, which leads to the notion of *partial termination*, a concept useful for reasoning about process state. To explain partial termination, let us define a *list* as a pair (*head*, *tail*), where *head* is any

value and *tail* is a *list* value or *null*. An input stream is simply a list with an unbound tail. If processing an input stream with an unbound tail, the program may never terminate. However, if the input stream stops growing, then based on the single rule for a dataflow variable (see above), the program will eventually stop executing while it waits for the input stream to grow (*i.e.*, for a new tail to be bound). This is an example of partial termination. It has not terminated completely because if the input stream grows, the program will execute further, up to the next partial termination. If there are no further input, then the program will execute no further. Partial termination is particularly useful because it allows us to determine that a long running process is waiting on an external message.

The new notion of *partial activation* occurs when a new semantic stack is added to a partially terminated process, therefore guaranteeing at least one semantic reduction. The process is then partially activated, which means that it will execute, but it may or may not execute to complete termination as described above. Generalizing partial activation allows for a variety of activation scenarios in addition to the scenario above that uses a stream. Partial activation combined with partial termination allows us to define the life-cycle model for long-running processes. The life-cycle for long-running processes behaves as follows:

1) The process is created by initializing an abstract machine with an initial semantic stack that is reducible. The process is now partially activated and execution begins.
2) The process waits on an external message and cannot execute further because all semantic stacks are waiting on dataflow variables and therefore cannot be reduced. The process is now partially terminated and eligible to be saved to disk.
3) An external message arrives for a partially terminated process. The process is restored, from disk if necessary, and a new semantic stack is added to the process. The process is now partially activated and execution begins.
4) The process has been fully reduced, all semantic stacks are empty. The process is now completely terminated.

*Lazy execution* and *failed values* are two other OCM constructs important to services orchestration. Lazy execution allows for late binding of modules and supports advanced functional programming techniques. Late binding is crucial to large process configurations so that services are bound only when necessary. Failed values are necessary to properly communicate failures. When the computation of a dataflow variable fails, a failed value is stored in the dataflow variable and the failure is propagated to other threads waiting on that dataflow variable.

In summary, so far we have reviewed the Oz Kernel Language and Oz Abstract Machine, as well as the highlights of the theory behind the Oz Computational Model that are most relevant to services orchestration: implicit synchronization, dataflow variables, lazy execution and failed values. These constructs combined with the new notion of partial activation form the technical foundation for the Oz language interpreter presented later in this paper.

## IV. Ozy Architecture

The Ozy architecture is based on the general view of the network-based architecture (net-arch) [4]. The conceptual elements of the net-arch are *Components*, *Connectors* and *Data*. As illustrated in Fig. 4, components interact by sending data through communications connectors (message passing). Components, and only components, perform computation on state, which can be a combination of data passed on connectors and private data known only to the component. The connector is especially useful because it encapsulates the technology used to communicate with other services. Components, connectors and data messages are the basic abstractions used in the Ozy architecture.

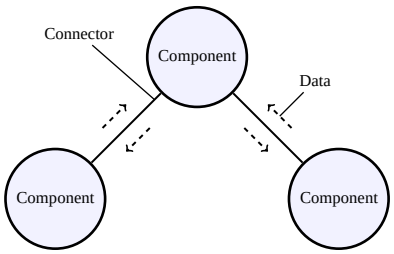


Figure 4. Elements of the Network-Based Architecture (net-arch) wherein components interact by sending data through communications connectors.

The manner in which Ozy employs the connector abstraction is what differentiates it from SOA-based and RESTful-based architectures. The SOA architecture and its BPEL are directly dependent and therefore constrained by the underlying specifications for WSDL and UDDI. The RESTful architecture is directly dependent on the uniform HTTP operations of GET, PUT, POST and DELETE. In contrast, the Ozy architecture loosens the constraints on the connector abstraction. A connector simply transfers data messages from one interface to another without any semantic transformation. Internally, the connector may be a subsystem that transforms and adapts, but when the message is delivered, it is semantically unchanged.

We must emphasize that the Ozy architecture is designed for hiding services implementations. Ozy does not define a new kind of services interface. It is very important that public interfaces for composite services be founded on industry standards, which is currently either SOAP or RESTful style interfaces. By moving the dependencies on network interfaces into connector abstractions (adapters), Ozy allows services containers to operate freely and evolve independently as long as service invocations can be adapted

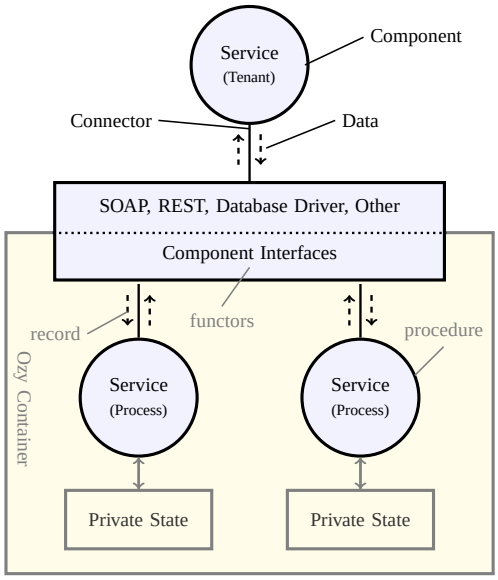


Figure 5. The Ozy Architecture. Ozy does not define a new services interface; industry standard SOAP or RESTful interfaces should still be employed, but the Ozy Architecture does provide a better separation of concern between interface and implementation.

between private and public interfaces. Furthermore, private-only compositions can be optimized to simply use local interfaces.

Ozy exploits the conceptual elements of the net-arch to create a better separation of concern between service interfaces and service implementations. As shown in Fig. 5, an Ozy interface can be SOAP or RESTful as long as the supporting connectors are available to adapt SOAP or RESTful messages to Oz language elements, such as functors, records and procedures. Likewise, the Ozy implementation can invoke any service; public or private, local or remote, given that the supporting connectors can translate Oz language elements into target interfaces. Finally, non-standard connectors can be installed to transparently support interactions with non-standard protocols, such as NoSQL drivers and legacy ERP interfaces.

## V. Ozy Container

The Ozy container is implemented in accordance with the Ozy architecture. The container hosts two fundamental activities: message routing and computation. The connector abstraction is realized by the computations that are performed along the message paths. When a message reaches a process interface, Ozy invokes an Oz language interpreter to compute the message.

Ozy employs an actor model for non-blocking asynchronous message routing. There are two types of messages: *Tell* and *Ask*. Unlike Tell, the Ask message requires a response and is built on an asynchronous Future/Promise library that routes responses back to the requester.

Message routing is intrinsically multi-tenant. Figure 6 illustrates the message routing in Ozy, where the

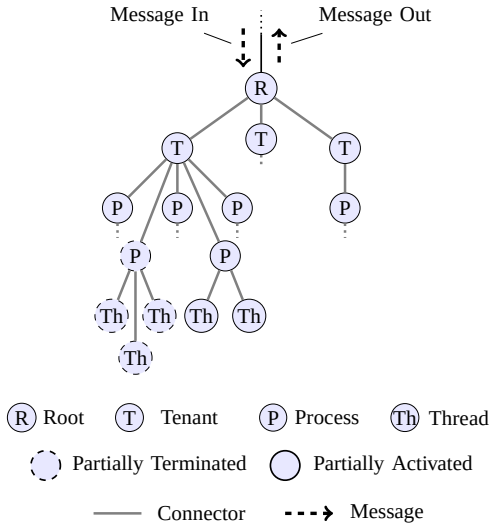


Figure 6. Message routing in Ozy. Message routing is based on the Ozy Architecture where the net-arch components are Root, Tenant, Process and Thread.

tenant and process services from Fig. 5 corresponds to the tenants and processes in Fig 6. Message routing is based on the following hierarchy template: `/root/tenants/{tenantId}/processes/{processId}` where curly braces denote identifying parameters. Messages are first received in the root actor mailbox where they are subsequently processed and routed to the appropriate tenant actor. The tenant actor receives request messages in its mailbox where they are subsequently processed and routed to process actors. If necessary, a tenant actor will activate a process actor prior to sending it a message if that process has been partially terminated and had its execution state persisted to disk.

Messages received from external services must be correlated to process instances and procedure invocations. A persistent correlation table is used to map business attributes to process identifiers. For example, suppose that a request to create a supply-order is being processed and a correlation entry is made that associates the supply-order number with the running process id. Now let us suppose the process partially terminates because the supply-order must be approved before it can be fulfilled. Later, when the supply manager approves (or disapproves) the supply-order we can use the correlation table to find the partially terminated process, activate it and allow it to continue executing.

```
1  proc next_process_message(p, m)
2      (s, E) = correlate_with_semantic_stack(p, m)
3      p.MST = [(s, E)] ⊎ p.MST
4      p.run()
5  end
```

Figure 7. Next Process Message. A message received by a Process must be correlated to a semantic statement. A typical correlation simply invokes a procedure.

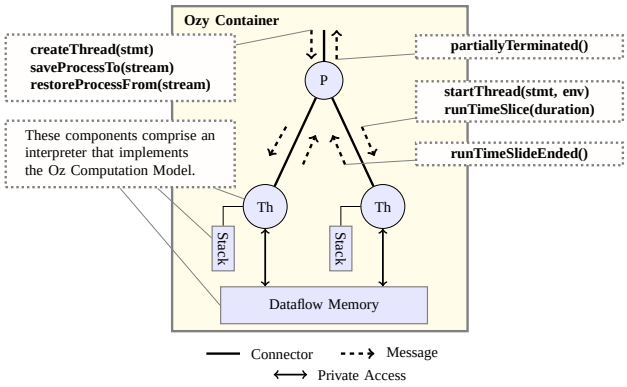


Figure 8. Ozy Computation. If we zoom-in on Process and Thread from Fig. 6, we see the intersection of message routing and computation. Process and Thread implement the Oz Computation Model.

Similarly, correlation entries must be made to map external attribute names to dataflow variables. Figure 7 shows how an incoming process message must be correlated to a semantic stack. Suppose that approving a supply-order involved setting an approved-flag to true or false. Now suppose that the `createSupplyOrder` process is partially terminated and waiting on the approved-flag dataflow variable known internally, and that a correlation entry exists with this fact. Later, when the supply manager approves (or disapproves) the supply-order, the approved-flag correlation entry can be used to construct a semantic stack that will set the approved-flag dataflow variable and allow the process to continue executing.

Figure 8 shows the intersection of message routing and computation. After a message is correlated and translated into a semantic stack, semantic computation can begin. The first step of computation is performed when Ozy selects a time slice message for execution, thereby satisfying the OCM definition of thread selection: MST = $\{ST\} \uplus MST'$. The computation step is completed when a thread actor runs a time slice by popping its semantic stack and performing a semantic operation as a function of its environment and the single-assignment store (dataflow memory), thereby satisfying the OCM definition of a full computation step: $(\{ST\} \uplus MST', \sigma) \rightarrow (\{ST'\} \uplus MST', \sigma')$.

## VI. Case Study

In this section we present two scenarios programmed with Ozy, which highlight the limitations of current approaches.

Our first scenario in Fig. 9 comes from the original JOLIE paper [20]. This scenario is a typical business process involving five participants: a customer, a market, a service register, a supplier and a bank. Like JOLIE, this example demonstrates an approach more expressive and more programmer-friendly than using BPEL XML; however, unlike JOLIE, this example uses implicit synchronization on dataflow variables and pattern matching, which are not

```
1  proc {GetPrice Quantity ClientLoc Product ?Euro} SupId ↩
       in
2      {Orch.updateCSet cset(product:Product quantity:↩
           Quantity clientLoc:ClientLoc)}
3      SupId = {Reg.getIdByQuery Product}
4      (SupLoc1 SupLoc2) = {Reg.getData SupId}
5      MyLoc = 'socket://localhost:2564'
6      Euro = {Supp.getEuro SupLoc1 Product} * Quantity
7      thread {Sleep 3000 millis} BuyTimeout=true end
8      case {WaitTwo BuyOk BuyTimeout}
9          of 1 then skip
10         [] 2 then {Quantity ClientLoc Product 'No'}
11     end
12 end
13 proc {Buy Quantity ClientLoc Product Conf} IdOrder BkId↩
       in
14     BuyOk = true
15     if Conf = 'Yes' then
16         IdOrder = {Supp.order SupLoc2 Quantity ↩
               ClientLoc Product}
17         BkId = {Bank.pay Reg.bankLoc ClientLoc SupLoc2 ↩
               MyLoc Euro}
18         {Orch.updateCSet cset(bkId:BkId)}
19         {Receipt BkId IdOrder}
20         {Orch.commit ClientLoc}
21     end
22 end
```

Figure 9. Typical Scenario. A typical business process that establishes a correlation set, finds services in the registry, and coordinates interactions with other business processes.

available in JOLIE, BPEL or other orchestration languages. Specifically, line 8 shows a pattern match implicitly synchronized on `BuyOk` and `BuyTimeout`. The pattern match waits on a nondeterministic choice from the `WaitTwo` procedure call. Computation is suspended at line 8 until either `BuyTimeout` is bound at line 7 or `BuyOk` is bound at line 14.

```
1  proc {Subscribe ClientLoc TankId ?Stream} Tail in
2      Stream = 0|Tail
3      {Orch.updateCSet cset(clientLoc:ClientLoc
4          tankId:TankId stream:Stream}
5      {PushNextEvent TankId Stream}
6  end
7  proc {PushNextEvent TankId Stream} in
8      thread {Sleep 1 minute} Proceed=true end
9      case {WaitTwo Stop Proceed}
10         of 1 then skip
11         [] 2 then
12             WaterLevel = {GetWaterLevel TankId}
13             if {WaterLevelChanged WaterLevel 0.05} Next↩
                   in
14                 case Stream of Head|Tail then
15                     Tail = WaterLevel|Next end
16             end
17             {PushNextEvent TankId Stream}
18     end
19 end
20 proc {Unsubscribe TankId Stream}
21     Stop=true
22 end
```

Figure 10. Non-typical Scenario. A non-typical Internet-of-Things scenario that emits an endless stream of integers to a waiting device until it unsubscribes.

Our second scenario in Fig. 10 is not a business process. Instead, it involves an Internet-of-Things device that monitors water level in a water tank. The service tracks the last water level reported and sends an update when the level changes by more than 5%. Water levels are gathered on the

local file system as comma-delimited text records. Line 2 initializes the `Stream` dataflow variable to a 0 water level value and an unbound tail. Line 3 establishes a correlation set with the `Stream` dataflow variable, which will be monitored by our connector. The procedure at Line 7 is called recursively (tail call optimized) from Line 17 to push water level events. Every time the `Stream` is extended, our connector is notified and transmits an integer to the device. This example demonstrates advanced programming features and flexibility not found in JOLIE, BPEL or other orchestration languages. Specifically, it uses a device connector based on a WebSocket where the connection is first established by the server for better security. Communication is a non-typical bi-directional connection that streams integers to the device only when needed at line 15, therefore avoiding inefficiencies and client-side polling. The service emits an endless stream of integers through the dataflow variable at line 15 until the connector unsubscribes at line 20.

## VII. Conclusions

Although BPEL XML and JOLIE have been used extensively to orchestrate business processes, they have major limitations that are succinctly yet powerfully addressed by applying lessons and technologies from the Oz computation model, most importantly implicit synchronization via dataflow variables and pattern matching. In this paper we present a description of the orchestration framework, Ozy, and provide its implementation as an open source project. We further show that the Ozy framework is easily used to program and perform standard business processes, and we also demonstrate the flexibility of Ozy with a long-running, implicitly synchronized, Internet-of-Things scenario.

This contribution opens the door for many possible follow on studies. One avenue for future work is in mobile devices and the Internet-of-Things, which are energy constrained. Ozy can be used for computational offloading, thereby reducing end-to-end latency between a device and its cloud services, which could be shown to improve the user experience and extend battery life.

## References


[1] Y. Duan, G. Fu, N. Zhou, X. Sun, N. C. Narendra, and B. Hu, "Everything as a service (xaas) on the cloud: Origins, current and future trends," in *Cloud Computing*. IEEE, 2015, pp. 621–628.

[2] M. P. Papazoglou, "Service-oriented computing: Concepts, characteristics and directions," in *WISE*. IEEE, 2003, pp. 3–12.

[3] A. Alves, A. Arkin, S. Askary, C. Barreto, B. Bloch, F. Curbera, M. Ford, Y. Goland, A. Guizar, N. Kartha *et al.*, "Web services business process execution language version 2.0," *Oasis Standard*. Oasis, 2007.

[4] R. T. Fielding, "Architectural styles and the design of network-based software architectures," Ph.D. dissertation, University of California, Irvine, 2000.

[5] M. Garriga, C. Mateos, A. Flores, A. Cechich, and A. Zunino, "Restful service composition at a glance: A survey," *Journal of Net. and Comp. App.*, vol. 60, pp. 32–53, 2016.

[6] R. Ranjan, R. Buyya, S. Nepal, and D. Georgakopulos, "A note on resource orchestration for cloud computing," *Concurrency and Computation: Practice and Experience*, vol. 27, no. 9, pp. 2370–2372, 2015.

[7] Gartner, "Market guide for rapid mobile app development," *Gartner Market Guide*, 2014.

[8] S. Yangui and S. Tata, "The spd approach to deploy service-based applications in the cloud," *Concurrency and Computation: Practice and Experience*, vol. 27, no. 15, pp. 3943–3960, 2015.

[9] P. Van-Roy and S. Haridi, *Concepts, techniques, and models of computer programming*. The MIT Press, 2004.

[10] G. W. Osborne and G. I. Mashini, "Method and system for simplifying object mapping for a user interface," Apr. 16 2013, US Patent 8,423,561.

[11] G. M. Medeiros Campos, N. Souto Rosa, and L. Ferreira Pires, "A survey of formalization approaches to service composition," in *Services Computing*. IEEE, 2014, pp. 179–186.

[12] C. Pautasso, "Restful web service composition with bpel for rest," *Data & Knowledge Engineering*, vol. 68, no. 9, pp. 851–866, 2009.

[13] P. Van Roy and S. Haridi, "Teaching programming with the kernel language approach," in *Workshop on Functional and Declarative Programming in Education*, 2002.

[14] G. Smolka, "The oz programming model," in *Computer science today*. Springer, 1995, pp. 324–343.

[15] M. Muller, T. Muller, and P. V. Roy, "Multiparadigm programming in oz," *Visions for the Future of Logic Programming: Laying the Foundations for a Modern Successor of Prolog*, 1995.

[16] L. Qilin and Z. Mintian, "The state of the art in middleware," in *IFITA*, vol. 1. IEEE, 2010, pp. 83–85.

[17] J. Al-Jaroodi and N. Mohamed, "Service-oriented middleware: a survey," *Journal of Network and Computer Applications*, vol. 35, no. 1, pp. 211–220, 2012.

[18] C. Guidi, R. Lucchi, R. Gorrieri, N. Busi, and G. Zavattaro, "Sock: a calculus for service oriented computing," in *Service Oriented Computing*. Springer, 2006, pp. 327–338.

[19] F. Montesi, C. Guidi, R. Lucchi, and G. Zavattaro, "Jolie: a java orchestration language interpreter engine," *Electronic Notes in Theoretical Comp. Sci.*, vol. 181, pp. 19–33, 2007.

[20] F. Montesi, C. Guidi, and G. Zavattaro, "Composing services with jolie," in *ECOWS*. IEEE, 2007, pp. 13–22.

[21] Q. Z. Sheng, X. Qiao, A. V. Vasilakos, C. Szabo, S. Bourne, and X. Xu, "Web services composition: A decades overview," *Information Sciences*, vol. 280, pp. 218–238, 2014.

[22] P. de Vrieze, L. Xu, A. Bouguettaya, J. Yang, and J. Chen, "Building enterprise mashups," *Future Generation Computer Systems*, vol. 27, no. 5, pp. 637–642, 2011.

[23] S. Balasubramaniam, G. Lewis, S. Simanta, D. B. Smith *et al.*, "Situated software: concepts, motivation, technology, and the future," *IEEE Software*, vol. 25, no. 6, pp. 50–55, 2008.

[24] C. Atkinson, P. Bostan, and D. Draheim, "A unified conceptual framework for service-oriented computing," in *Transactions on Large-Scale Data-and Knowledge-Centered Systems*. Springer, 2012, pp. 128–169.